TEASER

DDOS ATTACKS
AND PREVENTION TUTORIALS

LAYER 7: APPLICATION LEVEL DDOS

SUBVERTING BIND'S SRTT ALGORITHM
DERANDOMIZING NS SELECTION

CLOUDBASED DDOS PROTECTION SERVICES

ISSE 2014

14th & 15th October 2014
MCE, Brussels, Belgium
www.isse.eu.com

Securing Assets Across Europe

Europe's leading independent, interdisciplinary security conference and exhibition

Over the past decade, Information Security Solutions Europe (ISSE) has built an unrivalled reputation for its world-class, interdisciplinary approach and independent perspective on the e-security market.

This year, ISSE will take place on 14th & 15th October in Brussels. Regularly attracting over 300 professionals including government, commercial end-users and industry experts who will come together for a unique all-encompassing opportunity to learn, share and discuss the latest developments in e-security and identity management.

Programme Topic Areas

- **Trust Services, eID and Cloud Security**
  European trust services and eIdentity regulation, governance rules, standardization, interoperability of services and applications, architectures in the cloud, governance, risks, migration issues

- **BYOD and Mobile Security**
  Processes and technologies for managing BYOD programs, smartphone/tablet security, mobile malware, application threats

- **Cybersecurity, Cybercrime, Critical Infrastructures**
  Attacks & countermeasures against industrial Infrastructures; CERT/CSIRT – European & global developments, resilience of networks & services, surveillance techniques & analytics

- **Security Management, CISO Inside**
  CISOs featuring the latest trends and issues in information security, risk mitigation, compliance & governance; policy, planning and emerging areas of enterprise security architecture

- **Privacy, Data Protection, Human Factors**
  Issues in big data & cloud, privacy enhancing technologies, insider threats, social networking/engineering and security awareness programs

- **Regulation & Policies**
  Governmental cybersecurity strategies, authentication, authorization & accounting, governance, risk & compliance

In partnership with

eema
www.eema.org

TeleTrusT
Pioneers in IT security.

For more information visit www.isse.eu.com

@ISSEConference

DDoS Attacks and Prevention Tutorials

Table of Contents

Round Trip Time) algorithm. Our method enables derandomization of the target name server thus reduces the expected time of DNS cache poisoning attacks...

Layer 7: Application Level DdoS
by Neha Malik

Typically, a Denial of Service (DoS) condition occurs when a server or network resource is unable to service legitimate requests made to it, and therefore unable to perform a function it was designed to. DoS attacks have been around for some time, with the earliest attacks being dated to the first half of 1970's. This type of attack started out as an avenue for hackers to establish status in underground communities. However, these have evolved into far more sophisticated and dangerous forms that are directed at specific targets for a number of reasons, not excluding cyber-terrorism, corporate rivalry, hacktivism and even exhortation...

Tackling Layer 7 DDoS Attacks
by Ratan Jyoti

Distributed Denial of Service (DDoS) attack is a strewn challenge where the spurious or fake packets are sent to the victim in abnormally large number. DDoS attempts to block important services running on victim's server by flooding the victim's server with packets. The difference with DoS is that DDoS is that the attacks do not originate from a single host or network but from multiple hosts or networks which might have already been compromised. One of the main challenges here is to find the location of attackers and then block traffic at points near to the source of the attacks. Layer 7 DDoS attacks target the application layer at web or mail servers (eg. HTTP(S), SMTP, FTP etc) such that the service can be denied in effective way to bring web server to lock up or crash. Since they operate at the application protocol level which is OSI Layer 7, this attack is known as Layer 7 DDoS attack...

DDoS Attacks and Defense
by Rodrigo Salvalagio, Eder Plansky Silva

Denial of Service Attacks exists in decades, but frequency, extension and sophistication are evolving faster than companies can absorb them. This article shows both sides: how attackers are orchestrating and how companies are protecting themselves...

DdoS Attacks – Why and How?
by Satinder Sandhu

In the confidentiality, integrity, and availability metrics of information security, also known as the CIA Triad, denial-of-service (DoS) attacks impact availability. This article focuses on the types of attacks generally referred to as distributed denial-of-service, or DDoS, attacks. DDoS attacks are most often used to extort or damage businesses whose websites or online assets are a major source of revenue, are an indicator of brand value, or are critical to operations...

DDoS Attack – You Could Be Attacked Right Now. Are you prepared?
by Abdy Martinez

A DDoS attack could happen at any moment. Even if you have a powerful and well-configured firewall, updated anti-virus and anti-spam or good security practices, depending on the mode of DDoS attack, you will be affected. Trust me! So, are you ready to confront a DDoS attack? No? Please, before you start reading this article, ask your ISP a quote for DDoS protection service...

Protection Against DDoS On The Cloud
by Ahmed Fawzy

Simply the denial of service (Dos) attempts to deny legitimate users to the cloud service, in the (DDoS) the same matter happens but the attack lunched through thousands of zombies or fake packets may led to SLA violation, loses in revenue, lost productivity...

Dear Hakin9 Readers!

Let us present our latest issue entitled DdoS Attacks and Protection. Inside, you will find a few interesting tutorials that help develop your skills. Our experts prepared 11 articles in which they aim to familiarize you with various attacks and defence techniques.

We hope you enjoy the issue.

Krzysztof Samborski
and Hakin9 team

DDoS and the Internet

by Antonio Ieranò

Around security and Internet there is always a lot of talk, from time to time different kind of attack or threats comes out to the public attention and are overexposed by media and vendor marketing.

One truth we should always remember is that there are a lot of different attacks and, from time to time, one or another rises or peak due to several circumstances: political, environmental, technological or economical, there will always be a turnover of different technologies misused to perform an attack on the net.

One classic form of attack that has been on the news for a while, but then actually never stopped to exist, is the so called Denial of Service.

What is a denial of service?

From Wikipedia:

> *In computing, a denial-of-service (DoS) or distributed denial-of-service (DDoS) attack is an attempt to make a machine or network resource unavailable to its intended users. Although the means to carry out, motives for, and targets of a DoS attack may vary, it generally consists of efforts to temporarily or indefinitely interrupt or suspend services of a host connected to the Internet. As a clarification, DDoS (Distributed Denial of Service) attacks are sent by two or more persons, or bots. (See botnet) DoS (Denial of Service) attacks are sent by one person or system.*

In other words, a Denial of Service is a kind of attack that purpose is to stop, slow-down or somehow damage a service provided by someone in the network.

Denial-of-service attacks are considered violations of the Internet Architecture Board's Internet proper use policy, and violate the acceptable use policies of virtually all Internet service providers. They also commonly constitute violations of the laws of individual nations.

A DDoS (distributed denial of service) or a Dos (denial of service) are not related to specific technologies or process. Any hacking technique could be used to DDoS or dos a target, and sometimes we can have the same result using simple legal technique, or just because of misconfigurations.

To have a better understanding of what is a ddos attack we should first of all start from the very basic:

A dos attack aim is to damage (stop, slowdown..) a service

So as a first instance we should have a service running somewhere and someone else willing to attack it.

I will not take in account configuration errors or hardware crash since we are trying to understand the voluntary will to stop someone else service.

The situation is more or less the following:

Services running

Figure 1. Service running and the hacker

We have a service running somewhere and a hacker trying to stop (or….) the service.

In order to be able to perform this operation the hacker can target different areas of the process that allow the service to run:

Services running exposures

Figure 2. Service running exposures

Some of those attacking surface are common to any service running others can be service specific. So we will always have an Operating System running while a Database running is a need only for specific services.

- Every attacking area can be target to perform a Dos attack, so, as an example to stop a ecommerce service we could run different operations:

- We could try to hack the OS, and, as an example, escalate the administration rights in order to stop the service itself

- We could use a vulnerability related to the service itself to make it crash with, again a simple and understandable example, a memory buffer overflow.

- We could otherwise target a related service that is essential to the main objective function, think of targeting the database that is holding the information running on a wordpress site.

- But we can simply try to saturate the IP bandwidth of the server running the service in order to stop it,

- Or we can try to modify the network path of the service itself

And so on.

The result will always be the same; the service will stop or slow down.

In terms of what I can do to obtain the service denial I can:

- Try to saturate the resources related to that service

- Try to operate on the configuration parameters of something related to the service.

Resource saturation (Starvation Attacks)

When I try to perform a resource saturation, I need to exceed the computational or network resources of a specific sub target.

This kind of attack is quite common, and usually require that more point of attack works together in order to saturate the target providing a number of excessive request and or specifically crafted ones that overload the processing capability.

Figure 3. Resource Saturation

Normal target for that kind of operation are CPU, Disk and Network.

It is interesting to notice that a resource saturation can be effective even if performed on the target not directly related to the service we want to stop.

This means that the exposed area for a Dos is extremely Wide.

We can start considering what happen when we want to saturate the HW resources: as disk I\O or CPU.

One simple way to obtain this kind of result is just overloading the service with requests. Any request that need an answer will use CPU and disk resources, and since resources are limited the amount of instances that can be processed have a physical limit.

If we are able to exceed this limit we will be able to perform a system stop, at least till the system will be able to answer to all requests in the queue.

This simple DDoS attack have a great advantage, since it can use legit requests in order to obtain the needed hack.

This is typically the case of DDoS (Distribuited denial of service) which structure is, at its basic, very simple:

An attacker somehow takes control of a set of computer that are used to perform the attack. This is usually the botnet environment.

Figure 4. DDoS attack, the attacker use a multitude of computer to target the service

We should be aware that is particularly difficult to understand if we are in the presence of a DDoS attack or simply we have reached the limit of our configuration since this attack can be performed simply using a standard service request, the same request used by a usual user of the service itself.

We can understand we are in the presence of an attack when:

• It comes in a specific timeframe not related with normal operation

• It is a one-time event

• The request comes from unusual geographical location

• There are not systems that are exceeding resource consumption due to some scheduled or exceptional activity

• …

In other words, we should be aware of all the resources and their usual baseline, the origin of the traffic we receive and the kind of resources used by the surrounding services.

A DDoS can be performed not only by Botnet, but is also a typical acktivism [1] manifestation, several campaigns or Anonymous and the other groups are able to move a great number of members that can just start visiting a specific site all together blocking it.

It is easier to understand we are in presence of a DDoS when the kind of traffic is unusual or forget\ manipulated. As an example, an excess of ping can be a clear symptom of an incumbent DDoS attack.

In literature, we can find thousands of different kinds of attack that can provoke a DoS, most of them are network related, but the always present SQL injection is another classic example.

Typical Network attack include:

• Internet Control Message Protocol (ICMP) flood

• (S)SYN flood

• Teardrop attacks

• Peer-to-peer attacks

• Nuke

• …

But we can find the same kind of attack (flood, spoofing) ad any layer and even against the service itself.

Think of a service that have to process your request, if your request is complicated enough could bring it to halt or slow so from extremely complex regular expressions, to handling http request for an enormous amount of time anything ca be used.

Sometimes a simple SQL injection request or a simple buffer overflow are enough to consume all disk I\O or CPU resource

A particularly easy way to perform a DDoS attack is, last but not least, not to attack the service itself but to prevent users to reach the service.

 In this situation, the attacker wants to isolate the server providing the service targeting a key element in the chain that is used to reach the service itself.

Besides the obvious target of router or switch [2], a classical victim is the Name Server Structure. Attacking a DNS is quite easy and most effective due to the fact that most of the DNS running on the internet are poorly protected, heavily exposed and bad managed. This is a classical "Achilles' Heel" ☺.

Configuration attack

Another way to perform a Dos Attack is to conveniently hack the host and escalate credentials in order to take control of the system. Most of the time it is of no need to be the administrator to perform those kinds of attack, some sort of a power user are enough to modify any specific flag that can cause the damage. Of course, that kind of attack requires a specific set of hacking knowledge in order to penetrate the platform and escalate the credentials conveniently.

Bug, Backdoors, security holes everything can be used.

When we think about ddos related to configuration modification we should extend our analysis to all the network surrounding.

Taking control of a router or a switch can be extremely useful to perform a DDoS attack, modifying Routing Map or QoS configuration, altering ACL are all techniques that can be used to obtain the result.

Why, Who, What and When?

Why?

One question can rise, why anyone should perform a DDoS attack?

There are several reason to perform a DDoS or a Dos attack, some are evident some others can be more tricky.

The most evident is that someone want to stop a specific service, this can be done, at least, for 3 reasons:

• Political o Activist reasons, to have visibility or to demonstrate a specific idea.

• Retaliation, blocking a service and asking money to restore the functionality.

• Negative marketing, to give a bad feeling of the service provided (they do not run…)

But a DDoS attack can be part of something more complex, as an APT. typically in those situations the Dos attack is used to

• Mislead attention to something that is not the real target

• Cover track of activities

• As a needed part of the APT because the reaction at the attack will let the attacker to penetrate the system.

Both be the final activity o part of a more complex attack the Dos result effective targeting a quite various set of areas. To give an indication on what could be the target for a dos to a service is a hard exercise that require a complete analysis of the network and all the interaction between the service involved and the surrounding. Not all the attack need to be performed in the perimeter where the service reside, in a DNS attack the target could be the provider or even the root for that specific domain. As a result, a single technology that can help us against dos and ddos does not exist, but a set of technologies and procedures that runs form firewalls and IPS to dedicated DDoS technologies (that usually target a portion of the network exposure area) and may be some reputation services to analyze the origin of the request.

Who?

We sometimes misunderstood the extension and the deep of the DDoS phenomenon. One of the main reason is that there is not a correct perception of what a DDoS can do and who can cause it.

Due to the extreme variety of DDoS technicality and effect there are plenty of subjects that can perform a DDoS attack including people without specific knowledge, since DDoS tools are available on the internet, and some (think of the Low Orbit Ion Cannon) of public domain and easily available.

In the end any traffic generator or stress test tool can be used to perform a basic DDoS attack, but there are also sales Kit on the internet (the dark one, of course) and even "DDoS as a Service."

So it is not only a matter of hacker or activism (we all heard about Joker, Lulz, Anonymous, Iranian cyber army), but also criminality and even governments use this kind of techniques, the stuxnet affair was related to government (Israel? USA? Both?) vs. government (Iran) using a DDoS weaponized software (scada controller) through malware (stuxnet).

What?

Anything can be the target of a DDoS, anything that provide a service on a network is a good subject, so don't think there are areas that are "secure" or "safe". The question is if this service is meaningful for someone that can be targeted or want to target someone else?

Even training could be a good reason, or just a demonstration to rise the level of awareness or just create a white rumor on the background to cover real intention.

When?

Any time is good, the evolving political, economic and technological situation create moment by moment, a world of "good" reason for a DDoS. So Do not ask yourself "if" but "when"

References
[1] Richard Stallman has stated that DoS is a form of 'Internet Street Protests'
[2] May be someone remember the old classical Broadcast Storms, not always related to miscunfiguration of switches….

About the Author

Antonio Ieranò is an IT professional, marketing specialist, and tech evangelist with over 16 years of experience serving as a community liaison, subject matter expert, and high-profile trainer for key technologies and solutions. Mr Ieranò's experience includes acting as the public face of Cisco security technologies; leading pan-European technical teams in development of new Cisco security products; and serving as a key public speaker and trainer on behalf of new high-tech products. His expertise spans IT development and implementation, marketing strategy, legal issues, and budget / financial management.

Cloud-Based DDoS Protection Services

by Azi Ronen*, Chief Strategy Officer, SecurityDAM

In recent years Distributed Denial of Service (DDoS) attacks have become a mainstream threat to businesses, governmental agencies and critical infrastructure worldwide. DDoS attacks have grown in complexity, volume and sophistication. In a recent survey 65 percent of IT security practitioners reported experiencing an average of three DDoS attacks in the past 12 months [1].

With an average downtime of 54 minutes per attack and the cost amounting to as much as $100,000 per minute – it would have been expected that organizations put into practice preventative measures to protect their networks and business. However, this is far from being the case.

Many organizations do not employ any DDoS protection at all. Others rely on ISP solutions or use on-premises equipment, which at best can deflect a single type of attack. However, such solutions fail to provide adequate protection against multi-level attacks, and they lack the expertise to handle new types of attacks.

To ensure business continuity and provide solid DDoS protection, a different, multi-layer approach is needed – and such approach presents a new service opportunity for providers of managed security services (MSSPs).

Distributed Denial of Service attacks can be broadly categorized into two types:

- Volumetric attacks flood the victim with high volume of packets or IP flows, consuming network equipment and bandwidth resources. Some examples include SYN flood attacks (high packet-per-second attacks), large UDP packet floods (bandwidth attacks), and ICMP floods.

- Application attacks, also known as "low and slow" attacks, directly attack applications, servers of specific services, exploiting implementation weaknesses and design flaws. Some examples include HTTP Get or Post flood attacks, DNS flood attacks and SSL flood attacks.


Figure 1. Radware Security Survey

As can be seen in Figure 1, 27% of the DDoS attack saturate the victim internet link, using a volumetric attack. Such attacks cannot be mitigated by any local device.

The Hybrid Approach for DDoS prevention Services

In order to get a full protection against all types of DDoS attacks, a multi-layer solution is required. The solution is composed of the following main components:

- CPE (Customer-premises equipment) is a detection and signaling device placed at the edge of the customer's data center. Constantly monitoring network traffic, the CPE learns the traffic patterns to establish a normal behavior baseline. It detects anomalies and DDoS attacks early on, mitigates application attacks (❶ in Figure 2 below) and alerts the MSSP when the attacks are too large and saturate the enterprise access link.

- Scrubbing Centers, a cloud-based facility, manned by an emergency response team to ensure the fastest analysis and resolution of new attack types. When the network is under volumetric DDOS attack, traffic is redirected (❷ in Figure 2 below) to the scrubbing center for attack mitigation. After filtering, clean traffic is passed back to its original destination using GRE tunnels (❸ in Figure 2 below). Attack data is collected and stored, enabling real-time monitoring and historical reporting.

- A Customer's Portal, usually a web-based portal that provides real-time insight into events, attack characteristics, post-attack reports and statistics to the customers of the service.

Figure 2. Network under volumetric DDOS attack

The Business Opportunity for MSSPs

Providing cloud-based DDoS protection services provides a unique business opportunity to MSSPs. Local solutions deployed by enterprises at the data center cannot handle volumetric attacks and require the use of on-demand, cloud service that will be able to mitigate high-volume attacks that sometimes reach a volume higher than 100 GBPS.

A recent report by Infonetics Research [2] concludes that the "global cloud and CPE managed security service market grew another 12% in 2012, to $13 billion. While the majority of security service revenue in 2012 came from CPE-based services, by 2017 CPE revenue is expected to dip to 50% of total revenue. Infonetics forecasts sales of cloud-based security services to grow 69% over the next 5 years. "Other security services," of which hosted DDoS services are a major and growing contributor, are anticipated to comprise over 20% of cloud-based security service revenue by 2017."

Figure 3. Customers' Portal – the Dashboard

It is the right time for MSSPs to join the growing business of DDoS prevention services. SecurityDAM was established to enable the service infrastructure for MSSPs in the shortest time to market, based on state-of-the-art DDoS prevention and service management technology, and team of security experts providing the best service to customers.

References

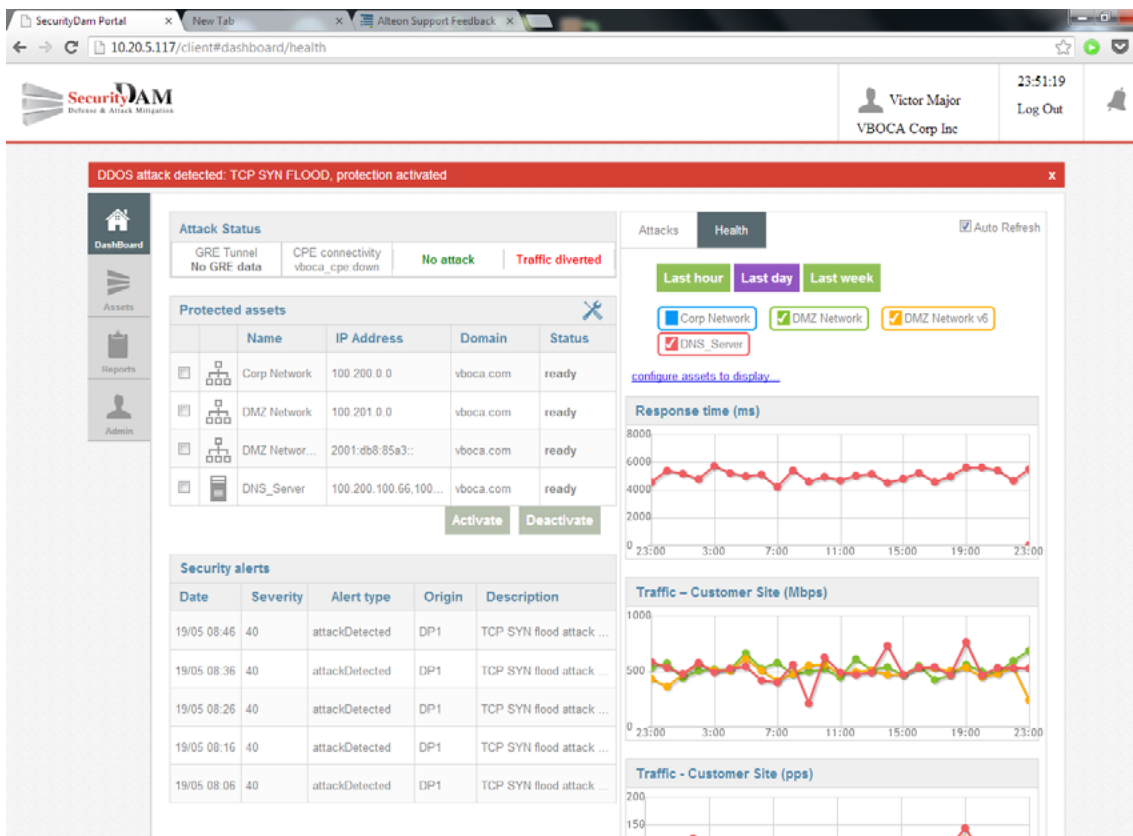[1] The research for Cyber Security on the Offense: A Study of IT Security Experts (*http://security.radware.com/uploadedFiles/Resources_and_Content/Attack_Tools/CyberSecurityontheOffense.pdf*), November 2012, by the Ponemon Institute and Radware.

[2] Managed security services top $13 billion in 2012 (*http://www.infonetics.com/pr/2013/2H12-Cloud-and-CPE-Managed-Security-Services-Market-Highlights.asp*); strong growth ahead for cloud security, April 2013, By Infonetics Research

About SecurityDAM

SecurityDAM provides world-class MSSP cloud-based solutions mitigating Distributed Denial of Service (DDoS) attacks on enterprise networks. Founded by a team of security experts, SecurityDAM is a member of the RAD group. For more information, see *www.securitydam.com*.

About the Author

*Azi Ronen has more than 20 years of marketing and product management experience in computer networking and fixed and mobile telecommunications, with strong technical background.
Prior to SecurityDAM, Mr. Ronen served as Executive Vice President of Marketing (6 years) and Corporate Development (4 years) at Allot Communications, a global leader in Broadband Traffic Management. This follows a career as VP Marketing of Vocaltec (VoIP space) and VP R&D/Product management of RADLINX, member of the RAD group.
Azi has a B.sc degree, Cum Laude, in Computer Sciences, from the Technion, Israel.

Layer 7: Application Level DDoS

by Neha Malik

How DDoS can be devastating for applications that aren't looking.

Typically, a Denial of Service (DoS) condition occurs when a server or network resource is unable to service legitimate requests made to it, and, therefore, unable to perform a function that it was designed to. DoS attacks have been around for quite some time, with the earliest attacks being dated to the first half of 1970's. This type of attack started out as an avenue for hackers to establish status in underground communities. However, today these have evolved into far sophisticated and dangerous forms that are directed at specific targets for a number of reasons, not excluding cyber-terrorism, corporate rivalry, hacktivism and even exhortation.

When a targeted Denial of Service attack is carried out using a large number of usually unwitting devices and internet connections across the world, it becomes a Distributed Denial of Service attack (DDoS). The skeletal structure of DDoS botnets is usually a variation of this:


Figure 1. Skeletal DDoS Anatomy

Distributed Denial of Service often thought of to be an attack concerned with Layer 3 or Layer 4 of the OSI Model. While DDoS defenses are getting better and more intelligent, attackers have managed to stay a step ahead all steps of the timeline. According to NSFOCUS, one major DDoS news event happened every 2 days, and one common DDoS attack happened every two minutes!

Despite this, attackers have not yet exploited the full range of vulnerabilities present in many online services for carrying out DDoS attacks. This is especially true of the Application and data processing layers. By nature, the application layer is more generic than the network layer due to a wide variety of applications. However, the methods to implement these applications are similar, leaving the application layer open to a large array of attacks, including unsophisticated ones. This explains these attacks are rapidly becoming the weapon of choice during recent years. According to a DDoS attack statistics analysis report by Prolexic, application level DDoS attacks consisted of 23.24% of total attacks in Q4 of 2013 alone. A number of Gartner reports show a disturbing deviation towards them as well.

Getting Popular

The motivation behind attackers going increasingly for application DDoS is simple:

Detectability

Layer 7 attacks are more difficult to detect than standard network-level DDoS attacks. The idea behind this assault is not the *appearance* of the data packets but the *intention.* Packets performing Layer 7 attacks look the same as any other legitimate packet to a firewall or IDS.

Efficiency

Unlike network attacks, application DDoS does not generate traffic spikes and alert detection mechanisms. Layer 7 DDoS is meaner and leaner in terms of lower consumption of bandwidth and requirement of intermediary resources.

Traceability

Application DDoS attacks use HTTP and HTTPS traffic. Traffic of malicious origin can be disguised via largely available proxy servers without much effort. Many proxy servers are notorious for not maintaining history or logs, making the attack much harder to trace back to its source.

Attack Anatomy

There are a number of attack vectors that have been used for exploitation of applications so far. Some of the most widely known are summarized below:

HTTP GET Flood Attack

This attack is carried out by bombarding the target server with a series of legitimate HTTP GET requests. This means that at this stage, the TCP three-way handshake has been completed and a valid connection has been established, deceiving Layer 4 detection devices. The idea behind the attack is to send a large number GET requests that are intended to exhaust server resources. Hence, attack vectors are made up resource-intensive requests like demands for large files or objects.

The logs for this type of attack look like any other request to the application:

```
80.93.170.6 – – [15/Apr/2014:19:40:08 -0530] "GET /?436873463892 HTTP/1.1" 200 440 "www.imdb.com/
   title/tt0298482" ""
179.10.10.92 – – [15/Apr/2014:19:40:08 -0530] "GET /?44328956742393 HTTP/1.1" 200 440 "www.
   youtube.com/playlist?list=PLF7A3E4527BCC3B56" ""
```

Particularly vulnerable are areas of search functionality within applications, especially those allow searches with wildcard characters, as these query databases and may lead to a database-level crash.

Some of the well-known tools for this type of attack include LOIC (Low Orbit Ion Canon), XOIC and HULK (HTTP Unbearable Load King).

HTTP POST Attack

This attack is executed via seemingly innocent requests with the *Content-Length* header manipulated to reflect a large value, e.g. 1000000. This tells the server how much content it should wait for before it

considers the request to be completed. Next, data is sent character by character over a very long period of time. The web server is forced to keep the connection open during this duration, affecting and/or denying legitimate user connections. This is especially fatal in a DDoS situation and may lead to the server crashing. It takes 20,000 such connections for an IIS server to be DDoS'ed!

Commonly used tools for this attack scenario are RUDY(R-U-Dead-Yet) and Tor's Hammer. OWASP has also come up with OWASP DoS HTTP POST tool.

HTTP Slow Read Attack

This attack works in a reverse way of HTTP POST attack. Instead of sending the server small requests, the approach taken is to reduce the client receive window to a very small size. Hence, server connections are compelled to stay open as the client reads responses indefinitely. This method bypasses server policies that filter slow-deciding customers. This attack proved to be successful when the following two conditions are satisfied:

The response size is large. This is easily satisfied with many web pages reaching sizes of up to 1MB.

Server send buffer size is known or can be estimated and the client receive buffer size is accordingly made small. The default value of send buffers is usually between 65Kb and 128Kb.

SlowHTTPTest can be used to carry out or test for HTTP Slow Read attack.

Slowloris Attack

Slowloris holds connections open by sending partial HTTP requests, particularly at the header. It works by sending incomplete header information distributed over long periods of time, thus holding up the server. Web servers look for a double carriage return to understand the end of a HTTP header. However, Slowloris continues sending information without providing the header's end.

The Slowloris tool is capable of modifying sent headers depending upon the target host configuration. For high traffic websites, the attacker may have to wait for all sockets to become available in order to consume the web server resources.

NTP Amplification Attack

2013 was the year of DNS Amplification. However, 2014 seems to be the year of NTP Amplification, with a reported rise of 371% in the first quarter. For a DNS Amplification attack, the amplification factor (ratio of response to request) is 8X. For an NTP attack on a busy server, it can reach 206X! An attacker, armed with the list of open NTP servers available on the internet, sends an NTP *monlist* (or `MON_GETLIST`) command with the source IP spoofed to be the target's IP address. The result is a very large response split over multiple packets directed to the target, leading to a Denial of Service condition.

Monlist modules can be found in Nmap as well as Metasploit.

SNMP Amplification Attack

Unlike other DDoS attacks, SNMP allows attackers to take over network devices as well and use them as bots in attacking other targets. To execute this attack, the attacker needs a list of exploitable SNMP hosts as well as community strings. This can be obtained by port-scanning IP addresses or obtaining the SNMP host list through private sources. In the next step, the attack sends a SNMP `BulkGetRequest` command to the SNMP Management Information Base (MIB), which returns amplified content. This attack request expectedly made with the source IP spoofed, leading to the target being overwhelmed with SNMP responses. Snmpbulkwalk is one of the tools used for this purpose.

So far, known instances of SNMP Amplification are comparatively lesser in number. This could be accounted to the lower visibility of SNMP servers over the Internet and the additional password requirements. However, the theoretical amplification factor of SNMP attacks has been found to be 650X – It is better to be safe than sorry here!

SMTP DDoS Attacks

SMTP DDoS can happen through several attack vectors. The first way is when thousands of emails are sent using computers across the web to one single SMTP server.

The second way is through backscattering attacks, where the attacker forces the SMTP server to generate a large number of non-delivery reports. Since non-delivery reports often include the full body of the original message along with attachments, the multiplicative force of this affect creates a DoS condition.

PyLoris tool is known to be used for execution of protocol-based, and specifically SMTP, DDoS attacks.

Application Logic Attacks

Apart from the various widespread attacks that are known to be launched against Layer 7, there is also a category of attacks that is seemingly overlooked and is not well-defined as of yet. These include application business logic and implementation logic flaws. Possible attack vectors can include the following:

- Exploitation of hidden bottlenecks in the application architecture. For example, applications that implement large client-facing tiers but have a small resource farm at the back-end to handle client requests.

- Applications implementing poor data validation techniques and thus, being DoS'ed by common injection flaws.

- Automated submission of data through Dictionary attacks for logins, overloading application functions, deliberately invoking race conditions or attacking multiple entities.

- User data manipulation leading to unexpected server errors or opening up known exploitable vulnerabilities.

- For applications locking out accounts permanently on entering invalid credentials, intentional account lockout of all accounts, including self, resulting in Denial of Service.

- Creation of multiple fake users for applications that do not require manual intervention at registration and thereby, starving application resources.

- Exhaustion of application session resources by creating an excessive number of active connections.

Shield of Defense

What we have seen previously are few known attacks that have been carried out against applications for creating DoS conditions. The actual number and type of assaults is continually growing and changing. Hence, it is imperative that organizations also evolve to keep up with the attackers. The best defense mechanisms are always first proactive, then reactive. Therefore, security measures for protection start at Design. In addition to specific precautions that are required for unique attacks, the following broad behaviors are a must for having a sound Defense in Depth structure.

Analysis of Logical Flaws

As the threat landscape for application attacks in general and DDoS in particular comes in focus, the golden rule of Input Validation becomes even more important for web application design. From an attacker's perspective, the first line of thought is almost always what an application user is allowed to

do. It is imperative that user input and actions not be implicitly trusted or assumed, and be thoroughly validated before consumption. Validation should be done at every application layer for the relevant layer, as input that is harmless for one layer may be intended for another. Secure Code Reviews and Penetration Tests hold special importance in this defense path.

Use of Anti-DDoS Tools and Testing

Many large vendors in the market have come up with DDoS solutions customized to include application protection. Some examples are solutions provided by Rackspace, F5, Juniper and others, including Akamai's KONA Site Defender, F5 BIG-IP Application Security Manage (ASM), Sucuri WAF, Cisco Traffic Anomaly Detector XT and Cisco Guard XT. Applications in question also need to undertake simulated attacks on a regular basis to understand their reaction to the same. While regular Penetration Testing is carried out for most Internet-facing application, it is crucial for these tests include examination for Denial of Service and Brute Force attacks.

Active Monitoring and Update

Finally, live monitoring mechanisms need to be in place to look out for unexpected application behavior. While a portion of this might be taken over by anti-DDoS solutions in place, it is necessary to have alert processes in place specifically for application attacks, along with SLA's defined for code changes by developers to stop attacks in question for good. Apart from these, regular and fast updates of vulnerable software as patches are released, is imperative. Other techniques like Blackholing requests can also be used. However, Blackholing discards legitimate traffic as well and may not be the best solution in an Application DDoS scenario.

It is important to understand that there is no such thing as "100% security" (Remember Titanic?). A combination of the above measures along with keeping updated on latest attack techniques can provide effective protection against Application DDoS in the long run.

On the Web
- *http://www.slideshare.net/prolexic7885/prolexic-d-do-s-attack-report-q4-2013-ddos-attack-trends-and-statistics* – DDoS attack trends report for Q4 2013
- *http://en.nsfocus.com/SecurityReport/2013%20NSFOCUS%20Mid-Year%20DDoS%20Threat%20Report.pdf* – 2013 Mid Year DDoS Threat Report
- *https://media.blackhat.com/bh-dc-11/Brennan/BlackHat_DC_2011_Brennan_Denial_Service-Slides.pdf* – Black Hat Presentation on DoS attacks
- *http://blog.cloudflare.com/understanding-and-mitigating-ntp-based-ddos-attacks* – Article on NTP Amplification attacks

About the Author

The author has over 5 years of experience in Information Security and currently works at KPMG Advisory Services. She was previously a Developer and Security Technologist at Microsoft Corporation.